

# Crossroads 2021 Technical Requirements Document

LA-UR-18-25993  
SAND2018-7366 O

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Triad National Security, LLC, for the National Nuclear Security Administration of the U.S. Department of Energy under contract 89233218CNA000001. LA-UR-18-25993. Approved for public release; distribution is unlimited.

Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND2018-7366 O.

# Crossroads 2021: Technical Requirements

---

<b>1</b>	<b>INTRODUCTION</b>	<b>4</b>
<b>1.1</b>	<b>SCHEDULE</b>	<b>6</b>
<b>2</b>	<b>SYSTEM DESCRIPTION</b>	<b>6</b>
<b>2.1</b>	<b>ARCHITECTURAL DESCRIPTION</b>	<b>6</b>
<b>2.2</b>	<b>SOFTWARE DESCRIPTION</b>	<b>7</b>
<b>2.3</b>	<b>PRODUCT ROADMAP DESCRIPTION</b>	<b>7</b>
<b>2.4</b>	<b>RISK MITIGATION STRATEGY</b>	<b>7</b>
<b>3</b>	<b>TARGETS FOR SYSTEM DESIGN, FEATURES, AND PERFORMANCE METRICS</b>	<b>7</b>
<b>3.1</b>	<b>SCALABILITY</b>	<b>8</b>
<b>3.2</b>	<b>SYSTEM SOFTWARE AND RUNTIME</b>	<b>10</b>
<b>3.3</b>	<b>SOFTWARE TOOLS AND PROGRAMMING ENVIRONMENT</b>	<b>12</b>
<b>3.4</b>	<b>PLATFORM STORAGE</b>	<b>15</b>
<b>3.5</b>	<b>APPLICATION PERFORMANCE</b>	<b>18</b>
<b>3.6</b>	<b>RESILIENCE, RELIABILITY, AND AVAILABILITY</b>	<b>22</b>
<b>3.7</b>	<b>APPLICATION TRANSITION SUPPORT AND EARLY ACCESS TO ACES TECHNOLOGIES</b>	<b>23</b>
<b>3.8</b>	<b>TARGET SYSTEM CONFIGURATION</b>	<b>24</b>
<b>3.9</b>	<b>SYSTEM OPERATIONS</b>	<b>25</b>
<b>3.10</b>	<b>POWER AND ENERGY</b>	<b>27</b>
<b>3.11</b>	<b>FACILITIES AND SITE INTEGRATION</b>	<b>29</b>
<b>4</b>	<b>OPTIONS</b>	<b>33</b>
<b>4.1</b>	<b>UPGRADES, EXPANSIONS AND ADDITIONS</b>	<b>33</b>
<b>4.2</b>	<b>EARLY ACCESS DEVELOPMENT SYSTEM</b>	<b>34</b>
<b>4.3</b>	<b>TEST SYSTEMS</b>	<b>35</b>
<b>4.4</b>	<b>ON SITE SYSTEM AND APPLICATION SOFTWARE ANALYSTS</b>	<b>35</b>
<b>4.5</b>	<b>DEINSTALLATION</b>	<b>35</b>
<b>4.6</b>	<b>MAINTENANCE AND SUPPORT</b>	<b>35</b>
<b>5</b>	<b>DELIVERY AND ACCEPTANCE</b>	<b>38</b>

<b>5.1</b>	<b>PRE-DELIVERY TESTING</b>	<b>38</b>
<b>5.2</b>	<b>SITE INTEGRATION AND POST-DELIVERY TESTING</b>	<b>38</b>
<b>5.3</b>	<b>ACCEPTANCE TESTING</b>	<b>39</b>
<b>6</b>	<b>RISK AND PROJECT MANAGEMENT</b>	<b>39</b>
<b>7</b>	<b>DOCUMENTATION AND TRAINING</b>	<b>40</b>
<b>7.1</b>	<b>DOCUMENTATION</b>	<b>40</b>
<b>7.2</b>	<b>TRAINING</b>	<b>40</b>
<b>8</b>	<b>REFERENCES</b>	<b>41</b>
	<b>APPENDIX A: SAMPLE ACCEPTANCE PLAN</b>	<b>42</b>
	<b>APPENDIX B: TRIAD SPECIFIC PROJECT MANAGEMENT REQUIREMENTS</b>	<b>50</b>
	<b>DEFINITIONS AND GLOSSARY</b>	<b>55</b>

# 1 Introduction

The Department of Energy (DOE) National Nuclear Security Administration (NNSA) Advanced Simulation and Computing (ASC) Program requires a computing system be deployed in 2021 to support the Stockpile Stewardship Program. In response to this requirement, Triad National Security, LLC (TNS), in furtherance of its participation in the Alliance for Computing at Extreme Scale (ACES), a collaboration between Los Alamos National Laboratory and Sandia National Laboratories, is releasing a Request for Proposal (RFP) for a next generation system, Crossroads.

In the 2021 timeframe, Trinity, the first ASC Advanced Technology System (ATS-1), will be nearing the end of its useful lifetime. Crossroads, the proposed ATS-3 system, provides a replacement, tri-lab computing resource for existing simulation codes and provides a resource for ever-increasing computing requirements to support the weapons program. The Crossroads system, to be sited at Los Alamos, NM, is projected to provide a large portion of the ATS resources for the NNSA ASC tri-lab simulation community: Los Alamos National Laboratory (LANL), Sandia National Laboratories (SNL), and Lawrence Livermore National Laboratory (LLNL), during the 2021-2026 timeframe. Crossroads is required to support stockpile stewardship certification and assessments to ensure that the nation's nuclear stockpile is safe, reliable and secure.

The ASC Program is faced with significant challenges resulting from the ongoing technology revolution. The program must continue to meet mission needs while adapting to sometimes radical changes in technology. Codes running on NNSA Advanced Technology Systems (Trinity and Sierra) in the 2019 timeframe are expected to run efficiently on Crossroads.

The goal of the Crossroads platform procurement is Efficiency. Efficiency will be evaluated in the areas of:

- Porting efficiency
- Performance efficiency
- Workflow efficiency

*Throughout this document, the term efficiency will refer to efficiency in these three areas unless otherwise specified.*

Trinity (ATS-1) will be used as the baseline for evaluating these goals. Porting efficiency is defined as the ease in which NNSA mission codes can be ported to execute on the proposed architecture. Minimal change to the existing code base is of high value. Performance efficiency is defined as the achieved performance of the application once ported to the proposed platform. Workflow efficiency is defined as the efficiency that a complete

NNSA workflow executes on the proposed platform. When evaluating proposals efficiency in all three stated areas will be considered together.

For example, a poor result would be a scenario where an application requires little porting effort to execute on the proposed platform but the resulting performance of the application is poor compared to the baseline. Ideally, individual applications that comprise a workflow can be easily ported to the proposed platform and perform well when compared with the baseline system. If, however, a necessary service like IO or efficient scheduling of a required resource for the workflow is inferior and hampers overall workflow efficiency this would still be a poor result.

To help inform the Offeror of the characteristics of NNSA workflows an accompanying whitepaper, "Crossroads Workflows," is provided that describes how application teams use High Performance Computing (HPC) resources today to advance scientific goals. The whitepaper is designed to provide a framework for reasoning about the optimal solution to these challenges. (The workflows document can be found on the Crossroads website <http://crossroads.lanl.gov/>.)

An Offeror's Technical Proposal shall include narrative and graphics, as appropriate, providing its responses/proposed solutions to each of the numbered sections of this Technical Requirements Document. An Offeror shall incorporate its responses/proposed solutions directly into each of the numbered sections of the Technical Requirements Document. The Technical Requirements Document is provided in MS Word format to facilitate this proposal requirement.

The evaluation committee will make no presumption of technical capability when evaluating an Offeror's responses/proposed solutions to this Technical Requirements Document and may downgrade a proposal if the Offeror's responses/proposed solutions are not materially responsive.

Where the word "should" appears throughout this document, it is used to convey a target that an Offeror ought to meet or exceed. If an Offeror exceeds a target, its proposal will be upgraded. If an Offeror fails to meet a target, its proposal will be downgraded.

Where the word "shall" appears throughout this document, it is used to impose a requirement that an Offeror must meet or exceed. If an Offeror fails to meet a requirement, its proposal will be downgraded or deemed non-responsive.

Each response/proposed solution shall clearly describe the role of any lower-tier subcontractor(s) and the technology or technologies, both hardware and software, and value added that the lower-tier subcontractor(s) provide, where appropriate.

The scope of work and technical specifications for any subcontracts resulting from this RFP will be negotiated based on this Technical Requirements Document and the successful Offeror's responses/proposed solutions.

Crossroads has a maximum funding limit over the system lifetime, to include all design and development, site preparation, maintenance, support and analysts. Total Cost of Ownership (TCO) will be considered in system selection. The Offeror must respond with configuration and pricing for both the primary and alternate point designs.

## 1.1 Schedule

The following is the tentative schedule for the Crossroads system.

*Table 1 Crossroads Schedule*

RFP Released	Q1CY19
On-site System Delivery Begins	Q2CY21
On-site System Delivery Complete	Q3CY21
Acceptance Complete	Q1CY22

## 2 System Description

### 2.1 Architectural Description

The Offeror shall provide a detailed full system architectural description of the Crossroads systems, including diagrams and text describing the following details as they pertain to the Offeror's system architectures (primary and alternate):

- Component architecture – details of all processor(s), memory technologies, storage technologies, network interconnect(s) and any other applicable components.
- Node architecture(s) – details of how components are combined into the node architecture(s). Details shall include bandwidth and latency specifications (or projections) between components.
- Board and/or blade architecture(s) – details of how the node architecture(s) is integrated at the board and/or blade level. Details should include all inter-node and inter-board/blade communication paths and any additional board/blade level components.
- Rack and/or cabinet architecture(s) – details of how board and/or blades are organized and integrated into racks and/or cabinets. Details should include all inter rack/cabinet communication paths and any additional rack/cabinet level components.

- Platform storage – details of how storage is integrated with the system, including a platform storage architectural diagram.
- System architecture – details of how rack or cabinets are combined to produce system architecture, including the high-speed interconnects and network topologies (if multiple) and platform storage.
- Proposed floor plan – including details of the physical footprint of the system and all of the supporting components, including details of site and facility integration requirements (e.g. power, cooling, and network).

## **2.2 Software Description**

The Offeror shall provide a detailed description of the proposed software eco-system, including a high-level software architectural diagram. Specify the provenance of the software component, for example open source or proprietary, and support mechanism for each (for the lifetime of the system including updates).

## **2.3 Product Roadmap Description**

The Offeror shall describe how the system does or does not fit into the Offeror's long-term product roadmap and a potential follow-on system acquisition in the 2025/26 and beyond timeframe.

## **2.4 Risk Mitigation Strategy**

The Offeror shall provide a summary of an alternate risk mitigation point design. The alternate point design shall be based on an architecture that reduces the risk of successful on-time deployment, for example, poses less schedule risk for delivery. It is of great importance that a viable platform (primary or alternate) is delivered in the Crossroads timeframe capable of supporting mission needs regardless of unforeseen technology disruptions. The Offeror shall not submit a full alternative point design proposal. Instead its summary of the alternate risk mitigation point design shall clearly describe any differences from the primary design point and how each of the noted differences satisfy the technical requirements contained in this document and reduces schedule risk for delivery of Crossroads.

# **3 Targets for System Design, Features, and Performance Metrics**

This section contains targets for detailed system design, features and performance metrics. It is desirable that the Offeror's proposal meet or exceed the targets outlined in this section. If a target cannot be met, the Offeror shall provide a development and deployment plan, including a schedule, to satisfy the target.

The Offeror may also propose any hardware and/or software architectural features that will provide improvements for any aspect of the system.

### **3.1 Scalability**

The scale of the system necessary to meet the needs of the application performance, porting and workflow requirements of the NNSA laboratories adds significant challenges. The Offeror should propose a system that enables efficiency at up to the full scale of the system. Additionally, the system proposed should provide functionality that assists users in enhancing efficiency at up to full scale. Scalability features, both hardware and software, that benefit both current and future programming models are essential. Memory bandwidth and latency are often limiting factors in the performance of NNSA mission applications therefore high value will be put on features that increase memory bandwidth or lower memory latency.

- 3.1.1 The system should support running jobs up to and including the full scale of the system.
- 3.1.2 The system should support launching an application at full system scale in less than 30 seconds. The Offeror shall describe factors (such as executable size) that could potentially affect application launch time.
- 3.1.3 The Offeror shall describe how application launch scales with the number of concurrent launch requests (per second) and the scale of each launch request (resources requested, such as the number of schedulable units etc.), including information such as:
  - All system-level and node-level overhead in the process startup including how overhead scales with node count for parallel applications, or how overhead scales with the application count for large numbers of serial applications.
  - Any limitations for processes on compute nodes from interfacing with an external work-flow manager, external database or message queue system.
- 3.1.4 The system should support at least 1000 concurrent users and more than 20,000 concurrent batch jobs. The system should allow a single user to execute multiple independent applications on a subset or all of the pool of nodes allocated to them. The Offeror shall describe details, including limitations of their proposed support for this requirement.
- 3.1.5 The Offeror shall describe all areas of the system in which node-level resource usage (hardware and software) increases as a job scales up (node, core or thread count).



- 3.1.6 The system should utilize an optimized job placement algorithm to reduce job runtime, lower variability, minimize latency, etc. The Offeror shall describe in detail how the algorithm is optimized to the system architecture.
- 3.1.7 The system should include an application programming interface to allow applications access to the physical-to-logical mapping information of the job's node allocation – including a mapping between MPI ranks and network topology coordinates, and core, node and rack identifiers.
- 3.1.8 The system software solution should provide a low jitter environment for applications and should provide an estimate of a compute node operating system's noise profile, both while idle and while running a non-trivial MPI application. If core specialization is used, the Offeror shall describe the system software activity that remains on the application cores.
- 3.1.9 The system should provide correct numerical results and consistent runtimes (i.e. wall clock time) that do not vary more than 3% from run to run in dedicated mode and 5% in production mode. The Offeror shall describe strategies for minimizing runtime variability.
- 3.1.10 The system's high speed interconnect should support a high messaging bandwidth, high injection rate, low latency, high throughput, and independent progress. The Offeror shall describe:
- The system interconnect in detail, including any mechanisms for adapting to heavy loads or inoperable links, as well as a description of how different types of failures will be addressed.
  - How the interface will allow all cores in the system to simultaneously communicate synchronously or asynchronously with the high speed interconnect.
  - How the interconnect will enable low-latency communication for one- and two-sided paradigms.
- 3.1.11 The Offeror shall describe how both hardware and software components of the interconnect support effective computation and communication overlap for both point-to-point operations and collective operations (i.e., the ability of the interconnect subsystem to progress outstanding communication requests in the background of the main computation thread).
- 3.1.12 The Offeror shall report or project the proposed system's node injection/ejection bandwidth.
- 3.1.13 The Offeror shall report or project the proposed system's bit error rate of the interconnect in terms of time period between errors that interrupt a job running at the full scale of the system.

3.1.14 The Offeror shall describe how the interconnect of the system will provide Quality of Service (QoS) capabilities (e.g., in the form of virtual channels or other sub-system QoS capabilities), including but not limited to:

- An explanation of how these capabilities can be used to prevent core communication traffic from interfering with other classes of communication, such as debugging and performance tools or with I/O traffic.
- An explanation of how these capabilities allow efficient adaptive routing as well as a capability to prevent traffic from different applications interfering with each other (either through QoS capabilities or appropriate job partitioning).
- An explanation of any sub-system QoS capabilities (e.g. platform storage QoS features).

3.1.15 The Offeror shall describe specialized hardware or software features of the system that enhance workflows or components of workflow efficiency, and describe any limits to their scalability on the system. The hardware should be on the same high speed network as the main compute resources and should have equal access to other compute resources (e.g. file systems and platform storage). It is desirable that the hardware have the same node level architecture as the main compute resources, but could, for example, have more memory per node.

## **3.2 System Software and Runtime**

The system should include a well-integrated and supported system software environment. The overall imperative is to provide users with a productive, high-performing, reliable, and scalable system software environment that enables efficient use of the full capability of the system.

- 3.2.1 The system should include a full-featured Linux operating system environment on all user visible service partitions (e.g., front-end nodes, service nodes, I/O nodes). The Offeror shall describe the proposed full-featured Linux operating system environment.
- 3.2.2 The system should include an optimized compute partition operating system that provides an efficient execution environment for applications running up to full-system scale. The Offeror shall describe any HPC relevant optimizations made to the compute partition operating system.
- 3.2.3 The Offeror shall describe the security capabilities of all operating systems proposed, e.g. compute, service.

- 3.2.4 The system should support a cohesive and integrated solution for launching user applications both at scale and high request frequency that require data at runtime such as: shared objects, containerized objects, data files, and dependent software.
- 3.2.5 The system should include resource management functionality, including job migration, backfill, targeting of specified resources (e.g., platform storage, CPU, memory), advance and persistent reservations, job preemption, job accounting, architecture-aware job placement, power management, job dependencies (e.g., workload management), resilience management, and high throughput workload execution (e.g., 100,000 job submissions per night). The Offeror may propose multiple solutions for a vendor-supported resource manager and should describe the benefits of each.
- 3.2.6 The system should support jobs consisting of multiple individual applications running simultaneously (inter-node or intra-node) and cooperating as part of an overall multi-component application (e.g., a job that couples a simulation application to an analysis application). The Offeror shall describe in detail how this will be supported by the system software infrastructure (e.g., user interfaces, security model, and inter-application communication).
- 3.2.7 The system should include a mechanism that will allow users to provide containerized software images without requiring privileged access to the system or allowing a user to escalate privilege. The startup time for launching a parallel application in a containerized software image at full system scale should not greatly exceed the startup time for launching a parallel application in the vendor-provided image.
- 3.2.8 The system should include a mechanism for dynamically configuring external IPv4/IPv6 connectivity to and from compute nodes, enabling special connectivity paths for subsets of nodes on a per-batch-job basis, and allowing fully routable interactions with external services.
- 3.2.9 The Successful Offeror should provide access to source code, and necessary build environment, for all software except for firmware, compilers, and third party products. The Successful Offeror should provide updates of source code, and any necessary build environment, for all software over the life of the subcontract.
- 3.2.10 The scheduler should support job workflows with data stage-in and stage-out from local file systems and storage systems accessible only from a remote data transfer system.

### 3.3 Software Tools and Programming Environment

The primary programming models used in production applications in this time frame are the Message Passing Interface (MPI), for inter-node communication, and OpenMP, for fine-grained on-node parallelism. While MPI+OpenMP will be the majority of the workload, the ACES laboratories expect some new applications to exercise emerging asynchronous programming models. System support that would accelerate these programming models/runtimes and benefit MPI+OpenMP is desirable.

- 3.3.1 The system should include an implementation of the most current version of MPI standard specification. The Offeror shall provide a detailed description of the MPI implementation (including specification version) and support for features such as hardware accelerated collective operations. The Offeror shall describe any limitations relative to the MPI standard.
- 3.3.2 The Offeror shall describe at what parallel granularity the system can be utilized by MPI-only applications.
- 3.3.3 The system should include optimized implementations of collective operations utilizing both inter-node and intra-node features where appropriate, including MPI\_Barrier, MPI\_Allreduce, MPI\_Reduce, MPI\_Allgather, and MPI\_Gather.
- 3.3.4 The Offeror shall describe the network transport layer of the system including any support for OpenUCX, Portals, libfabric, libverbs, and any other transport layer, including any optimizations of their implementation that will benefit application performance or workflow efficiency.
- 3.3.5 The system should include a complete implementation of the most current version of OpenMP standard including, if applicable, accelerator directives, as well as a supporting programming environment. The Offeror shall provide a detailed feature description of the OpenMP implementation(s) and describe any expected deviations from the OpenMP standard.
- 3.3.6 The Offeror shall provide a description of how applications written to utilize OpenMP will be compiled and executed on the system.
- 3.3.7 The Offeror shall provide a description of any proposed hardware or software features that enable OpenMP performance optimizations.
- 3.3.8 The Offeror shall list any PGAS languages and/or libraries that are supported (e.g. UPC, SHMEM/OpenSHMEM, CAF, Global Arrays) and describe any hardware and/or programming environment software that optimizes any of the listed PGAS languages supported on the system. The Offeror shall describe interoperability with MPI+OpenMP.

- 3.3.9 The Offeror shall describe and list support for any emerging programming models such as asynchronous task/data models (e.g., Legion, STAPL, HPX, or OCR) and describe any system hardware and/or programming environment software it will provide that optimizes any of the supported models. The Offeror shall describe interoperability with MPI+OpenMP.
- 3.3.10 The Offeror shall describe the proposed hardware and software environment support for:
- Fast thread synchronization of subsets of execution threads;
  - Atomic add, fetch-and-add, multiply, bitwise operations, and compare-and-swap operations over 32-bit and 64-bit integers, single-precision, and double-precision operands;
  - Atomic compare-and-swap operations over 16-byte wide operands that comprise two double precision values or two 64-bit memory pointer operands;
  - Fast context switching or task-switching;
  - Fast task spawning for unique and identical task with data dependencies;
  - Support for active messages.
- 3.3.11 The Offeror shall describe in detail all programming APIs, languages, compilers and compiler extensions, etc. other than MPI and OpenMP (e.g. OpenACC, CUDA, OpenCL, etc.) that will be supported by the system. It is desirable that instances of all programming models provided be interoperable and efficient when used within a single process or single job running on the same compute node.
- 3.3.12 The system should include support for the languages C, C++ (including complete support for C++11/14/17), Fortran 77, Fortran 90, and Fortran 2008 programming languages. Providing multiple compilation environments is highly desirable. The Offeror shall describe any limitations that can be expected in meeting full C++17 support based on current expectations. Key ASC applications push the limits of current Fortran compilers. The Offeror shall describe their support for Fortran, including standards levels and/or coverage of Fortran test suites, such as the [FLANG Fortran Test Suite](#).
- 3.3.13 The system should include a Python implementation that will run on the compute partition with optimized MPI4Py, NumPy, and SciPy libraries.
- 3.3.14 The system should include a programming toolchain(s) that enables runtime coexistence of threading in C, C++, and Fortran, from within applications and any supporting libraries using the same compiler toolchain. The Offeror shall describe the interaction between OpenMP and native parallelism expressed in language standards.

- 3.3.15 The system should include C++ compiler(s) that can successfully build the latest [Boost C++ library](#). The Offeror shall support the most recent stable version of Boost.
- 3.3.16 The system should include optimized versions of libm, libgsl, BLAS levels 1, 2 and 3, LAPACK, ScaLAPACK, HDF5, NetCDF, and FFTW. It is desirable for these to efficiently interoperate with applications that utilize OpenMP. The Offeror shall describe all other optimized libraries that will be supported, including a description of the interoperability of these libraries with the programming environments proposed.
- 3.3.17 The system should include a mechanism that enables control of task and memory placement within a node for efficient performance. The Offeror shall provide a detailed description of controls provided and any limitations that may exist.
- 3.3.18 The system should include a comprehensive software development environment with configuration and source code management tools. On heterogeneous systems, a mechanism (e.g., an upgraded autoconf) should be provided to create configure scripts to build cross-compiled applications on login nodes.
- 3.3.19 The system should include an interactive parallel debugger with an X11-based graphical user interface. The debugger should provide a single point of control that can debug applications in all supported languages using all granularities of parallelism (e.g. MPI+X) and programming environments provided and scale up to 25% of the system.
- 3.3.20 The system should include a suite of tools for detailed performance analysis and profiling of user applications. At least one tool should support all granularities of parallelism in mixed MPI+OpenMP programs and any additional programming models supported on the system. The tool suite must provide the ability to support multi-node integrated profiling of on-node parallelism and communication performance analysis. The Offeror shall describe all proposed tools and the scalability limitations of each. The Offeror shall describe tools for measuring I/O behavior of user applications.
- 3.3.21 The system should include event-tracing tools. Event tracing of interest includes: message-passing event tracing, I/O event tracing, floating point exception tracing, and message-passing profiling. The event-tracing tool API should provide functions to activate and deactivate event monitoring multiple times during execution from within a process.
- 3.3.22 The system should include single- and multi-node stack-tracing tools. The tool set should include a source-level stack trace back, including an API that allows a running process or thread to query its current stack trace.

- 3.3.23 The system should include tools to assist the programmer in introducing limited levels of parallelism and data structure refactoring to codes using any proposed programming models and languages. Tool(s) should additionally be provided to assist application developers in the design and placement of the data structures with the goal of optimizing data movement/placement for the classes of memory proposed in the system.
- 3.3.24 The system shall support licensing for the programming environment (compilers, debuggers, optimization tools, optimized math libraries, etc.) for up to twenty (20) concurrent users at job sizes that span from 100's of small-scale jobs for a single user all the way up to a single job occupying the full scale (100% of the compute partition) of the platform.

### **3.4 Platform Storage**

Platform storage is certain to be one of the advanced technology areas included in any system delivered in this timeframe. The ACES laboratories anticipate these emerging technologies will enable new usage models. With this in mind, an accompanying whitepaper, "APEX Workflows," is provided that describes how application teams use HPC resources today to advance scientific goals. The whitepaper is designed to provide a framework for reasoning about the optimal solution to these challenges. The whitepaper is intended to help an Offeror develop a platform storage architecture response that accelerates the science workflows while minimizing the total number of platform storage tiers. The workflows document can be found on the Crossroads [website](#).

- 3.4.1 The system should include platform storage capable of retaining all application input, output, and working data for 12 weeks (84 days), estimated at a minimum of 12% of baseline system memory per day.
- 3.4.2 The system should include platform storage with a warranted durability or a maintenance plan such that the platform storage is capable of absorbing approximately two times the systems baseline memory per day for a nominal 5 years.
- 3.4.3 The Offeror shall describe how the system provides sufficient bandwidth to support a JMTTI/Delta-Ckpt ratio of greater than 200. See Table 2 Target System Configuration.
- 3.4.4 The Offeror shall describe how the storage system provides sufficient performance to asynchronously migrate 80% of memory (i.e. a checkpoint from 3.4.3) from the fastest tier to the capacity tier in 75% of JMTTI.

- 3.4.5 The Offeror shall describe how the system satisfies a minimum storage bandwidth requirement capable of writing 25% of baseline system memory in less than 300 seconds.
- 3.4.6 The Offeror shall describe how a job running across the entire system with a MPI rank per core can create a file from every MPI rank in fewer than 10 seconds between the first and last create. If the response requires more than a single pre-existing directory the offeror shall also describe the required directory layout and the time required to create those directories.
- 3.4.7 The Offeror shall describe all available interfaces to platform storage for the system, including but not limited to:
- POSIX
  - APIs
  - Exceptions to POSIX compliance.
  - Time to consistency and any potential delays for reliable data consumption.
  - Any special requirements for users to achieve performance and/or consistent data.
- 3.4.8 The Offeror shall describe the reliability characteristics of platform storage, including but not limited to:
- Any single point of failure for all proposed platform storage tiers (note any component failure that will lead to temporary or permanent loss of data availability).
  - Enumerate platform storage tiers that are designed to be less reliable or do not use data protection techniques (e.g., replication, erasure coding).
  - Describe the impacts to a running compute job due to storage-related failures and during the recovery from said failure for each reliable platform tier. Specifically describe the job impact during failure, and separately describe the job impact during recovery.
  - Vendor supplied mechanisms to ensure data integrity for each platform storage tier (e.g., data scrubbing processes, background checksum verification, etc.).
  - Login or interactive nodes access to platform storage when the compute nodes are unavailable.
- 3.4.9 The Offeror shall describe system features for platform storage tier management designed to accelerate workflows, including but not limited to:



- Mechanisms for migrating data between platform storage tiers, including manual, scheduled, and/or automatic data migration to include rebalancing, draining, or rewriting data across devices within a tier.
- How platform storage will be instantiated with each job if it needs to be, and how platform storage may be persisted across jobs.
- The capabilities provided to define per-user policies and data movement between different tiers of platform storage or external storage resources (e.g., archives).
- Describe any data-related consistency, whether optional or inherent, between storage tiers (e.g. write-back caching).
- The ability to integrate with a scheduling resource.
- Mechanism to incrementally add capacity and bandwidth to a particular tier of platform storage. Please also describe functional and performance impacts to running jobs while the system integrates new resources.
- Capabilities to manage or interface platform storage with external storage resources or archives (e.g., fast storage layers or HPSS).

3.4.10 The Offeror shall describe software features that allow users to optimize I/O for the workflows of the system, including but not limited to:

- Batch data movement capabilities, especially when data resides on multiple tiers of platform storage.
- Methods for users to create and manage platform storage allocations.
- Any ability to directly target a tier for writing or reading data.
- Locality-aware job/data scheduling.
- Methods for users to exploit any enhanced performance of relaxed consistency.
- Methods for enabling user-defined metadata with the platform storage solution.

3.4.11 The Offeror shall describe the method and rate for enumerating the entire platform storage metadata. Describe any special capabilities that would mitigate user performance issues and/or allow the enumeration to complete in fewer than 4 hours; expect at least 1 billion objects.

3.4.12 The Offeror shall describe capabilities to comprehensively collect platform storage usage data and note those that can be collected out-of-band. Storage metrics for the system may include, but are not limited to:

- Per client metrics and frequency of collection, including but not limited to: the number of bytes read or written, number of read or write invocations, client cache statistics, and metadata statistics such as number of opens, closes, creates, and other system calls of relevance to the performance of platform storage.
- Job level metrics, such as the number of sessions each job initiates with each platform storage tier, session duration, total data transmitted (separated as reads and writes) during the session, and the number of total platform storage invocations made during the session.
- Platform storage tier metrics and frequency of collection, such as the number of bytes read, number of bytes written, number of read invocations, number of write invocations, bytes deleted/purged, number of I/O sessions established, and periods of outage/unavailability.
- Job level metrics describing usage of a tiered platform storage hierarchy, such as how long files are resident in each tier, hit rate of file pages in each tier (i.e., whether pages are actually read and how many times data is re-read), fraction of data moved between tiers because of a) explicit programmer control and b) transparent caching, and time interval between accesses to the same file (e.g., how long until an analysis program reads a simulation generated output file).

3.4.13 The Offeror shall propose a method for providing access to platform storage from other systems at the facility. In the case of tiered platform storage, at least one tier must satisfy this requirement.

3.4.14 The Offeror shall describe the capability for platform storage tiers to be repaired, serviced, and incrementally patched/upgraded while running different versions of software or firmware without requiring a storage tier-wide outage. The Offeror shall describe the level of performance degradation, if any, anticipated during the repair or service interval.

3.4.15 The Offeror shall specify the time required and the optimal number of compute nodes required to achieve peak read and write performance to the fastest platform storage tier using the following data sets:

- A 1 TB data set of 20 GB files.
- A 5 TB data set of any chosen file size. Offeror shall report the file size chosen.
- Usable capacity of the fastest tier using 32 MB files.

### **3.5 Application Performance**

Assuring that real applications perform efficiently on Crossroads is key for their success. Because the full applications are large, often with millions of

lines of code, and in some cases are export controlled, a suite of benchmarks have been developed for RFP response evaluation and system acceptance. The benchmark codes are representative of the workloads of the NNSA laboratories but often smaller than the full applications.

The performance of the benchmarks will be evaluated as part of both the RFP response and system acceptance. Final benchmark acceptance performance targets will be negotiated after a final system configuration is defined. All performance tests must continue to meet negotiated acceptance criteria throughout the lifetime of the system.

System acceptance for Crossroads will include export controlled ASC codes (to include code at 0D999 and ITAR control level) from each of the three NNSA laboratories. The benchmark information and licensing requirements regarding the Crossroads acceptance codes, and supplemental materials can be found on the Crossroads [website](#).

- 3.5.1 The Offeror shall provide responses for the benchmarks (SNAP, HPCG, PENNANT, MiniPIC, UMT, VPIC, Branson) provided on the *Crossroads benchmarks* link on the Crossroads [website](#). All modifications or new variants of the benchmarks (including makefiles, build scripts, and environment variables) are to be supplied in the Offeror's response.
- The results of all problem sizes (baseline and optimized) should be provided in the Offeror's Scalable System Improvement (SSI) spreadsheets. SSI is the calculation used for measuring improvement and is documented on the Crossroads [website](#), along with the SSI spreadsheets. If predicted or extrapolated results are provided, the methodology used to derive them should be clearly documented.
  - The Offeror shall provide licenses for the system for all compilers, libraries, and runtimes used to achieve benchmark performance.
- 3.5.2 The Offeror shall provide performance results for the system that may be benchmarked, predicted, and/or extrapolated for the baseline MPI+OpenMP variants of the benchmarks. The Offeror may modify the benchmarks to include extra OpenMP pragmas as required, but the benchmark must remain a standard-compliant program that maintains existing output subject to the validation criteria described in the benchmark run rules.
- 3.5.3 The Offeror shall optionally provide performance results from an Offeror optimized variant of the benchmarks. The Offeror may modify the benchmarks, including the algorithm and/or programming model used to demonstrate high system performance. If algorithmic changes are made, the Offeror shall provide an explanation of why the results may deviate from validation criteria described in the benchmark run rules.

- 3.5.4 In addition to the *Crossroads benchmarks*, an ASC Simulation Code Suite representing the three NNSA laboratories will be used to judge performance at time of acceptance (Mercury from Lawrence Livermore, PartiSN from Los Alamos, and SPARC from Sandia). NNSA mission requirements forecast the need for a 6X or greater improvement over the ASC Trinity system (Haswell partition) for the code suite, measured using SSI. Final acceptance performance targets will be established during negotiations after a final system configuration is defined. Information regarding the ASC Simulation Code Suite can be found on the Crossroads [website](#). Source code will be provided to the Offeror, but it will require compliance with export control laws and no cost licensing agreements.
- 3.5.5 The Offeror shall report or project the number of cores necessary to saturate the available node baseline memory bandwidth as measured by the Crossroads memory bandwidth benchmark found on the Crossroads [website](#).
- If the node contains heterogeneous cores, the Offeror shall report the number of cores of each architecture necessary to saturate the available baseline memory bandwidth.
  - If multiple tiers of memory are available, the Offeror shall report the above for every functional combination of core architecture and baseline or extended memory tier.
- 3.5.6 The Offeror shall report or project the sustained dense matrix multiplication performance on each type of processor core (individually and/or in parallel) of the system node architecture(s) as measured by the Crossroads multithreaded DGEMM benchmark found on the Crossroads [website](#).
- The Offeror shall describe the percentage of theoretical double-precision (64-bit) computational peak, which the benchmark GFLOP/s rate achieves for each type of compute core/unit in the response, and describe how this is calculated.
- 3.5.7 The Offeror shall report, or project, the MPI two-sided message rate of the nodes in the system under the following conditions measured by the communication benchmark specified on the Crossroads [website](#):
- Using one, two, four, eight, and half the number of cores of MPI ranks per node with MPI\_THREAD\_SINGLE.
  - Using one, two, four, eight, and half the number of cores of MPI ranks per node and multiple threads per rank with MPI\_THREAD\_MULTIPLE.
  - The Offeror may additionally choose to report on other configurations, including MPI\_THREAD\_SERIALIZED and MPI\_THREAD\_FUNNELED.

- 3.5.8 The Offeror shall report, or project, the MPI one-sided message rate of the nodes in the system for all passive synchronization RMA methods with both pre-allocated and dynamic memory windows under the following conditions measured by the communication benchmark specified on the Crossroads [website](#) using:
- One, two, four, eight, and half the number of cores of MPI ranks per node with MPI\_THREAD\_SINGLE.
  - One, two, four, eight, and half the number of cores of MPI ranks per node and multiple threads per rank with MPI\_THREAD\_MULTIPLE.
  - The Offeror may additionally choose to report on other configurations, including MPI\_THREAD\_SERIALIZED and MPI\_THREAD\_FUNNELED.
- 3.5.9 The Offeror shall report, or project, the time to perform the following collective operations for 25%, 50%, and 100% of the compute partition nodes in the system and report on core occupancy during the operations measured by the communication benchmark specified on the Crossroads [website](#) for:
- An 8 byte MPI\_Allreduce operation.
  - An 8 byte per rank MPI\_Allgather operation.
- 3.5.10 The Offeror shall report, or project, the minimum and maximum off-node latency of the system for MPI two-sided messages using the following threading modes measured by the communication benchmark specified on the Crossroads [website](#):
- MPI\_THREAD\_SINGLE with a single thread per rank.
  - MPI\_THREAD\_MULTIPLE with two or more threads per rank.
- 3.5.11 The Offeror shall report, or project, the minimum and maximum off-node latency for MPI one-sided messages of the system for all passive synchronization RMA methods with both pre-allocated and dynamic memory windows using the following threading modes measured by the communication benchmark specified on the Crossroads [website](#):
- MPI\_THREAD\_SINGLE with a single thread per rank.
  - MPI\_THREAD\_MULTIPLE with two or more threads per rank.
- 3.5.12 The Offeror shall provide an efficient implementation of MPI\_THREAD\_MULTIPLE. Bandwidth, latency, and message throughput measurements using the MPI\_THREAD\_MULTIPLE thread support level should have no more than a 10% performance degradation when compared to using the MPI\_THREAD\_SINGLE support level as measured by the communication benchmark specified on the Crossroads [website](#).

- 3.5.13 The Offeror shall report, or project, the maximum I/O bandwidths of the system as measured by the IOR benchmark specified on the Crossroads [website](#).
- 3.5.14 The Offeror shall report, or project, the metadata rates of the system as measured by the MDTEST benchmark specified on the Crossroads [website](#).
- 3.5.15 The Successful Offeror shall be required at time of acceptance to meet specified targets for acceptance benchmarks, and mission codes for Crossroads, listed on the Crossroads [website](#).
- 3.5.16 The Offeror shall describe how the system may be configured to support a high rate and bandwidth of TCP/IP connections to external services both from compute nodes and directly to and from the platform storage, including:
- Compute node external access should allow all nodes to each initiate 1 connection concurrently within a 1 second window.
  - Transfer of data over the external network to and from the compute nodes and platform storage at 100 GB/s per direction of a 1 TB dataset comprised of 20 GB files in 10 seconds.

### **3.6 Resilience, Reliability, and Availability**

The ability to achieve the NNSA mission goals hinges on the productivity of system users. System availability is therefore essential and requires system-wide focus to achieve a resilient, reliable, and available system. For each metric specified below, the Offeror must describe how they arrived at their estimates (e.g. failure rates of individual components including hardware and software that make up major aspects of Offeror's estimate).

- 3.6.1 Failure of the system management and/or RAS system(s) should not cause a system or job interrupt. This requirement does not apply to a RAS system feature, which automatically shuts down the system for safety reasons, such as an overheating condition.
- 3.6.2 The minimum System Mean Time Between Interrupt (SMTBI) should be greater than 720 hours.
- 3.6.3 The minimum Job Mean Time To Interrupt (JMTTI) should be greater than 24 hours. Automatic restarts do not mitigate a job interrupt for this metric.
- 3.6.4 The ratio of JMTTI/Delta-Ckpt should be greater than 200. This metric is a measure of the system's ability to make progress over a long period of time and corresponds to an efficiency of approximately 90%. If, for example, the JMTTI requirement is not met, the target JMTTI/Delta-Ckpt ratio ensures this minimum level of efficiency.

- 3.6.5 An immediate re-launch of an interrupted job should not require a complete resource reallocation. If a job is interrupted, there should be a mechanism that allows re-launch of the application using the same allocation of resource (e.g., compute nodes) that it had before the interrupt or an augmented allocation when part of the original allocation experiences a hard failure.
- 3.6.6 A complete system initialization should take no more than 30 minutes. The Offeror shall describe the full system initialization sequence and timings.
- 3.6.7 The system should achieve 99% scheduled system availability. System availability is defined in the glossary.
- 3.6.8 The Offeror shall describe the resilience, reliability, and availability mechanisms and capabilities of the system including, but not limited to:
- Any condition or event that can potentially cause a job interrupt.
  - Resiliency features to achieve the availability targets.
  - Single points of failure (hardware or software), and the potential effect on running applications and system availability.
  - How a job maintains its resource allocation and is able to relaunch an application after an interrupt.
  - A system-level mechanism to collect failure data for each kind of component.

### **3.7 Application Transition Support and Early Access to ACES Technologies**

The Crossroads system may include numerous advanced technologies. The Offeror shall include in their proposal a plan to effectively utilize these technologies and assist in transitioning the mission workflows to the system. The Successful Offeror shall support efforts to transition the Advanced Technology Development Mitigation (ATDM) codes to the systems. ATDM codes are currently being developed by the three NNSA weapons laboratories, Lawrence Livermore, Los Alamos, and Sandia. These codes may require compliance with export control laws and no cost licensing agreements. Information about the ATDM program can be found on the [NNSA website](#).

- 3.7.1 The Successful Offeror should provide a vehicle for supporting the successful demonstration of the application performance requirements and the transition of key applications to the Crossroads system (e.g., a Center of Excellence). Support should be provided by the Offeror and all of its key advanced technology providers (e.g., processor vendors, integrators, etc.). The Successful Offeror should provide experts in the areas of application porting and performance optimization in the form of staff training, general user training, and deep-dive interactions with a set of application code teams. Support should be provided from the date of subcontract execution through two (2) years after final acceptance of the systems.
- 3.7.2 The Successful Offeror shall describe their support structure for the proposed programming environment. This includes mechanisms for reporting issues and requesting new functionality, in addition to escalation paths/priorities available to Crossroads' applications. Support should be provided up to two (2) years after final acceptance of the systems.
- 3.7.3 The Offeror shall describe which of the proposed hardware and software technologies (physical hardware, emulators, and/or simulators) , will be available for access before system delivery and in what timeframe. The proposed technologies should provide value in advanced preparation for the delivery of the final Crossroads system for pre-system-delivery application porting and performance assessment activities.

### 3.8 Target System Configuration

ACES determined the following targets for Crossroads System Configurations. Offerors shall state projections for their proposed system configurations relative to these targets.

*Table 2 Target System Configuration*

Crossroads	
<b>Baseline Memory Capacity</b> <i>Excludes all levels of on-die-CPU cache</i>	<b>&gt; 0.5 PiB</b>
<b>Benchmark SSI increase over Trinity system (Haswell partition)</b>	<b>&gt; 6X</b>
<b>Platform Storage</b>	<b>&gt; 10X Baseline Memory</b>
<b>Nameplate Power</b>	<b>&lt; 20 MW</b>
<b>Peak Power</b>	<b>&lt; 18 MW</b>



Dated 07-19-18

<b>Nominal Power</b>	<b>&lt; 15 MW</b>
<b>Idle Power</b>	<b>&lt; 10% Nameplate Power</b>
<b>Job Mean Time to Interrupt (JMTTI)</b> <i>Calculated for a single job running on the entire system</i>	<b>&gt; 24 Hours</b>
<b>System Mean Time to Interrupt (SMTTI)</b>	<b>&gt;720 Hours</b>
<b>JMTTI/Delta-Ckpt</b>	<b>&gt; 200</b>
<b>System Availability</b>	<b>&gt; 99%</b>

### 3.9 System Operations

System management should be an integral feature of the overall system and should provide the ability to effectively manage system resources with high utilization and throughput under a workload with a wide range of concurrencies. The Successful Offeror should provide system administrators, security officers, and user-support personnel with productive and efficient system configuration management capabilities and an enhanced diagnostic environment.

- 3.9.1 The system should include scalable integrated system management capabilities that provide human interfaces and APIs for system configuration and its ability to be automated through configuration management software, software management, change management through a version control system, local site integration, and system configuration backup and recovery.
- 3.9.2 The system should include a means for tracking and analyzing all software updates, software and hardware failures, and hardware replacements over the lifetime of the system. All patches and releases should include changelogs with detailed descriptions of bug fixes and features and also what services are affected by these changes.
- 3.9.3 The system should include the ability to perform rolling upgrades and rollbacks on a subset of the system while the balance of the system remains in production operation. The Offeror shall describe the mechanisms, capabilities, workload management support, and limitations of rolling upgrades and rollbacks. No more than half the system partition should be required to be down for rolling upgrades and rollbacks.

- 3.9.4 The system should include an efficient mechanism for reconfiguring and rebooting compute nodes. The Offeror shall describe in detail the compute node reboot mechanism, differentiating types of boots (warmboot vs. coldboot) required for different node features, as well as how the time required to reboot scales with the number of nodes being rebooted. Warmboot timings for both file system and cluster boot shall be independently provided and also a combined warmboot file system and cluster boot timing if it differs.
- 3.9.5 The system should include a mechanism whereby all monitoring data and logs captured are available to the system owner, and will support an open monitoring API to facilitate lossless, scalable sampling and data collection for monitored data. Any filtering that may need to occur will be at the option of the system manager. The system will include a sampling and connection framework that allows the system manager to configure independent alternative parallel data streams to be directed off the system to site-configurable consumers.
- 3.9.6 The system should include a mechanism to collect and provide metrics and logs which monitor the status, health, utilization, and performance of the system, subsystems, and all major components, including, but not limited to:
- Environmental measurement capabilities for all systems and peripherals and their sub-systems and supporting infrastructure, including power and energy consumption and control.
  - Internal high speed network performance counters, including measures of network congestion and network resource consumption.
  - Information enabling traffic and congestion attribution, with explanation of the attribution logic.
  - Hardware performance counters enabling application performance assessment with the ability to integrate these with system metric (e.g., network performance counters) data.
  - All levels of integrated and attached platform storage.
  - The system as a whole, including hardware performance counters for metrics for all levels of integrated and attached platform storage.
- 3.9.7 The Offeror shall describe what tools and APIs it will provide for the collection, analysis, integration, and visualization of metrics and logs produced by the system (e.g., peripherals, integrated and attached platform storage, and environmental data, including power and energy consumption). The description should include any capabilities to configure collection rates for the available metrics.

- 3.9.8 The offeror shall describe all mechanisms for obtaining application performance and progress information such as enabling application specific logs and metrics to be collected and transported off the system or enabling hardware and system performance counters to be collected and transported off the system and associated with particular jobs/applications.
- 3.9.9 The Offeror shall describe the system configuration management and diagnostic capabilities of the system that address the following topics:
- Detailed description of the system management support.
  - Any effect or overhead of software management tool components on the CPU or memory available on compute nodes.
  - Release plan, with regression testing and validation for all system related software and security updates.
  - Support for multiple simultaneous or alternative system software configurations, including estimated time and effort required to install both a major and a minor system software update.
  - User activity tracking, such as audit logging and process accounting.
  - Unrestricted privileged access to all software and hardware components and all related performance metrics delivered with the system.
- 3.9.10 The system should provide a mechanism for reporting all basic component and essential services state (e.g., up/down/running) and changes of state. The system should also provide documented APIs for querying the state.
- 3.9.11 Offeror should provide a description of all fundamental data and associated metrics and computations used to assess status, health, utilization, and performance in 3.9.6.
- 3.9.12 The Offeror shall describe all measurement capabilities (system, rack/cabinet, board, node, component, and sub-component level) for the system, including control and response times, sampling frequency, accuracy of the data, and timestamps of the data for individual points of measurement and control.

### **3.10 Power and Energy**

Power, energy, and temperature will be critical factors in how the ACES laboratories manage systems in this time frame and must be an integral part of overall Systems Operations. The solution must be well integrated into other intersecting areas (e.g., facilities, resource management, runtime systems, and applications). The ACES laboratories expect a growing number of use cases in this area that will require a vertically integrated solution.

- 3.10.1 The Offeror shall describe all power, energy, and temperature operational measurement capabilities (system, rack/cabinet, board, node, component, and sub-component level) for the system, including control and response times, sampling frequency, accuracy of the data, and timestamps of the data for individual points of measurement and control.
- 3.10.2 The Offeror shall describe all operational control capabilities it will provide to affect power or energy use (system, rack/cabinet, board, node, component, and sub-component level).
- 3.10.3 The system should include system-level interfaces that enable measurement and dynamic control of power and energy relevant characteristics of the system, including but not limited to:
- AC measurement capabilities at the system or rack level.
  - System-level minimum and maximum power settings (e.g., power caps).
  - System-level power ramp up and down rate.
  - Scalable collection and retention all measurement data such as:
    - point-in-time power data.
    - energy usage information.
    - minimum and maximum power data.
- 3.10.4 The system should include resource manager interfaces that enable measurement and dynamic control of power and energy relevant characteristics of the system, including but not limited to:
- Job and node level minimum and maximum power settings.
  - Job and node level power ramp up and down rate.
  - Job and node level processor and/or core frequency control.
  - System and job level profiling and forecasting.
  - e.g., prediction of hourly power averages >24 hours in advance with a 1 MW tolerance.
- 3.10.5 The system should include application and runtime system interfaces that enable measurement and dynamic control of power and energy relevant characteristics of the system including but not limited to:
- Node level minimum and maximum power settings.
  - Node level processor and/or core frequency control.
  - Node level application hints, such as:

- application entering serial, parallel, computationally intense, I/O intense or communication intense phase.
- 3.10.6 The system should include an integrated API for all levels of measurement and control of power relevant characteristics of the system. It is preferable that the provided API complies with the High Performance Computing Power Application Programming Interface Specification (<http://powerapi.sandia.gov>).
- 3.10.7 The Offeror shall project (and report) the Nameplate, Peak, Nominal, and Idle Power of the system.
- 3.10.8 The Offeror shall describe any controls available to enforce or limit power usage below Nameplate power and the reaction time of this mechanism (e.g., what duration and magnitude can power usage exceed the imposed limits).
- 3.10.9 The Offeror shall describe the status of the system when in an Idle State (describe all Idle States if multiple are available) and the time to transition from the Idle State (or each Idle State if there are multiple) to the start of job execution.

### **3.11 Facilities and Site Integration**

- 3.11.1 The system should use 3-phase Delta 480V AC (four-wire system, three phases and one ground). Other system infrastructure components (e.g., disks, switches, login nodes, and mechanical subsystems such as CDUs) must use either 3-phase 480V AC (strongly preferred), 3-phase 208V AC (second choice), or single-phase 120/208V AC (third choice). The total number of individual branch circuits and phase load imbalance should be minimized.
- 3.11.2 All equipment and power control hardware of the system should be Nationally Recognized Testing Laboratories (NRTL) certified and bear appropriate NRTL labels.
- 3.11.3 Every rack, network switch, interconnect switch, node, and disk enclosure should be clearly labeled with a unique identifier visible from the front of the rack and the rear of the rack, as appropriate, when the rack door is open. These labels will be high quality so that they do not fall off, fade, disintegrate, or otherwise become unusable or unreadable during the lifetime of the system. Nodes will be labeled from the rear with a unique serial number for inventory tracking. It is desirable that motherboards also have a unique serial number for inventory tracking. Serial numbers shall be visible without having to disassemble the node, or they must be able to be queried from the system management console.

- 3.11.4 All components in a rack intended to be serviced while the rack has power shall be fully serviceable without danger of touching an exposed conducting surface. Consider power switches for individual components that may need to be powered-off/-on individually. Consider minimizing the number of connecting cables that need to be removed to power-off/-on a component. Consider the placement of connectors, handles, etc., with respect to conducting services, that must be used to remove and replace a component.
- 3.11.5 Table 3 below shows target facility requirements identified by ACES for the Crossroads system. The Offeror shall describe the features of its proposed systems relative to site integration at the respective facilities, including:
- Description of the physical packaging of the system, including dimensioned drawings of individual cabinets types and the floor layout of the entire system.
  - Remote environmental monitoring capabilities of the system and how it would integrate into facility monitoring.
  - Emergency shutdown capabilities.
  - Detailed descriptions of power and cooling distributions throughout the system, including power consumption for all subsystems.
  - Description of parasitic power losses within Offeror's equipment, such as fans, power supply conversion losses, power-factor effects, etc. For the computational and platform storage subsystems separately, give an estimate of the total power and parasitic power losses (whose difference should be power used by computational or platform storage components) at the minimum and maximum ITUE, which is defined as the ratio of total equipment power over power used by computational or platform storage components. Describe the conditions (e.g. "idle") at which the extrema occur.
  - OS distributions or other client requirements to support off-system access to the platform storage (e.g. LANL File Transfer Agents).

*Table 3 Crossroads Facility Requirements*

Location	Los Alamos National Laboratory, Los Alamos, NM. The system will be housed in the Strategic Computing Complex (SCC), Building 2327
Altitude	7,500 feet

Seismic	N/A
Water Cooling	<p>The system should operate in conformance with ASHRAE Class W2 guidelines (dated 2011). The facility will provide operating water temperature of 75°F, at up to 35PSI differential pressure at the system cabinets. However, Offeror should note if the system is capable of operating at higher temperatures.</p> <p>Note: LANL facility will provide inlet water at a nominal 75°F, per system design. Total flow requirements may not exceed 9600GPM.</p>
Water Chemistry	The system must operate with facility water meeting basic ASHRAE water chemistry. Special chemistry water is not available in the main building loop and would require a separate tertiary loop provided with the system. If tertiary loops are included in the system, the Offeror shall describe their operation and maintenance, including coolant chemistry, pressures, and flow controls. All coolant loops within the system should have reliable leak detection, temperature, and flow alarms, with automatic protection and notification mechanisms.
Air Cooling	The system must operate with supply air at 75°F-60°F, with a relative humidity from 30%-70%. The rate of airflow is between 800-1500 CFM/floor tile. No more than 3MW of heat should be removed by air cooling.
Maximum Power Rate of Change	The hourly average in system power should not exceed the 2MW wide power band negotiated at least 2 hours in advance.
Power Quality	The system should be resilient to incoming power fluctuations at least to the level guaranteed by the ITIC power quality curve.
Floor	42" raised floor
Ceiling	16-foot ceiling and 16-foot ceiling plenum

Maximum Footprint	8000 square feet; 80 feet long and 100 feet deep.
Shipment Dimensions and Weight	No restrictions.
Floor Loading	The average floor loading over the effective area should be no more than 300 pounds per square foot. The effective area is the actual loading area plus at most a foot of surrounding fully unloaded area. A maximum limit of 300 pounds per square foot also applies to all loads during installation. The Offeror shall describe how the weight will be distributed over the footprint of the rack (point loads, line loads, or evenly distributed over the entire footprint). A point load applied on a one square inch area should not exceed 1500 pounds. A dynamic load using a CISC Wheel 1 size should not exceed 1250 pounds (CISC Wheel 2 – 1000 pounds).
Cabling	All power cabling and water connections should be below the access floor. It is preferable that all other cabling (e.g., system interconnect) is above floor and integrated into the system cabinetry. Under floor cables (if unavoidable) should be plenum rated and comply with NEC 300.22 and NEC 645.5. All communications cables, wherever installed, should be source/destination labeled at both ends. All communications cables and fibers over 10 meters in length and installed under the floor should also have a unique serial number and dB loss data document (or equivalent) delivered at time of installation for each cable, if a method of measurement exists for cable type.
External network interfaces supported by the site for connectivity requirements specified below	1Gb, 10Gb, 40Gb, 100Gb, IB EDR, IB HDR. The network infrastructure is continuously upgraded moving to the latest Ethernet and IB capabilities.



External bandwidth on/off the system for general TCP/IP connectivity	Minimum of 100 GB/s per direction with a preference for 300 GB/s per direction. Describe how 100 GB/s per direction could be expanded to 300 GB/s per direction.
External bandwidth on/off the system for accessing the system's PFS	Minimum of 100 GB/s with a preference for 300 GB/s. Describe how 100 GB/s could be expanded to 300 GB/s.
External bandwidth on/off the system for accessing external, site supplied file systems. E.g. GPFS, NFS	Minimum of 100 GB/s with a preference for 300 GB/s. Describe how 100 GB/s could be expanded to 300 GB/s.

## 4 Options

The ACES team expects to have future requirements for system upgrades and/or additional quantities of components based on the configurations proposed in response to this solicitation. The Offeror should address any technical challenges foreseen with respect to scaling and any other production issues. Proposals should be as detailed as possible.

### 4.1 Upgrades, Expansions and Additions

The Offeror shall propose the following separately priced options using whatever is the natural unit for the proposed architecture design as determined by the Offeror. For example, for system size, the unit may be number of racks, number of blades, number of nodes or some other unit appropriate for the system architecture. If the proposed design has no option to scale one or more of these features, the Offeror should simply state this in the proposal response.

- 4.1.1 The Offeror shall describe and separately price options for scaling up the overall Crossroads system. These options may be larger than the smallest compute partition. Any of these options may be exercised multiple times.
  - 4.1.1.1 The Offeror shall propose a configuration or configurations which increase the baseline memory capacity in steps, e.g. by adding 25%, 50%, 100%, and 200%.

- 4.1.1.2 The Offeror shall propose and separately price upgrades or expansions for scaling the capacity and performance of the Crossroads I/O subsystem such that it can retain all application input, output, and working data for 24 and 36 weeks Refer to section where it indicates that the minimum amount of such data is 12% of baseline system memory per day.. If the Offeror's I/O subsystem consists of multiple storage tiers, the Offer shall describe and separately price options for scaling each storage tier separately.
- 4.1.1.3 The Offeror shall propose and separately price any options for upgrading the Offeror's proposed technology of the Crossroads system over its five-year lifetime.
- 4.1.2 The Offeror shall also provide separately priced options for systems which the CONTRACTOR may procure in addition to Crossroads system that provide approximately 10%, 25%, 50% and 200% of the Offeror's proposed capability of the Crossroads system. The options proposed for the expansion of the crossroads system under 4.1.1, above shall also apply to any additional system purchased at a cost proportional to the system capability compared to Crossroads.

## **4.2 Early Access Development System**

To allow for early and/or accelerated development of applications or development of functionality required as a part of the statement of work, the Offeror shall propose options for early access development systems. These systems can be in support of the baseline requirements or any proposed options.

- 4.2.1 The Offeror shall propose an Early Access Development System. The primary purpose is to expose the application to the same programming environment as will be found on the final system. It is acceptable for the early access system not to use the final processor, node, or high-speed interconnect architectures. However, the programming and runtime environment must be sufficiently similar that a port to the final system is trivial. The early access system shall contain similar functionality of the final system, including file systems, but scaled down to the appropriate configuration. The Offeror shall propose an option for the following configurations based on the size of the final Crossroads system.
- 2% of the compute partition.
  - 5% of the compute partition.
  - 10% of the compute partition.
- 4.2.2 The Offeror shall propose development test bed systems that will reduce risk and aid the development of any advanced functionality that is exercised as a part of the statement of work.

### **4.3 Test Systems**

The Offeror shall propose the following test systems. The systems shall contain all the functionality of the main system, including file systems, but scaled down to the appropriate configuration. Multiple test systems may be awarded.

- 4.3.1 The Offeror shall propose an Application Regression test system, which should contain at least 200 compute nodes.
- 4.3.2 The Offeror shall propose a System Development test system, which should contain at least 50 compute nodes.

### **4.4 On Site System and Application Software Analysts**

- 4.4.1 The Offeror shall propose and separately price two (2) System Software Analysts and two (2) Applications Software Analysts for each site. Offerors shall presume each analyst will be utilized for four (4) years. For Crossroads, these positions require a DOE Q-clearance for access.

### **4.5 Deinstallation**

The Offeror shall propose to deinstall, remove and/or recycle the system and supporting infrastructure at end of life. Storage media shall be destroyed to the satisfaction of ACES, and/or returned to ACES at its request.

### **4.6 Maintenance and Support**

The Offeror shall propose and separately price maintenance and support with the following features:

#### **4.6.1 Maintenance and Support Period**

The Offeror shall propose all maintenance and support for a period of four (4) years from the date of acceptance of the system. Warranty shall be included in the 4 years. For example, if the system is accepted on April 1, 2021 and the Warranty is for one year, then the Warranty ends on March 30, 2022, and the maintenance period begins April 1, 2022 and ends on March 30, 2025. Offeror shall also propose additional maintenance and support extension for years 5-7.

#### **4.6.2 Maintenance and Support Solutions**

The Offeror shall propose the following maintenance and support solutions and propose pricing separately for each solution. ACES may purchase either one of the solutions or neither of the solutions, at its discretion. Different maintenance solutions may be selected for the various test systems and final system.

#### 4.6.2.1 Solution 1 – 7x24

The Offeror shall price Solution 1 as full hardware and software support for all Offeror provided hardware components and software. The principal period of maintenance (PPM) shall be for 24 hours by 7 days a week with a four-hour response to any request for service. Hardware service requests require the Offeror to be on-site within four (4) hours of the request.

#### 4.6.2.2 Solution 2 – 5x9

The Offeror shall price Solution 2 as full hardware and software support for all Offeror provided hardware components and software. The principal period of maintenance (PPM) shall be on a 9 hours by 5 days a week (exclusive of holidays observed by ACES). The Successful Offeror shall provide hardware maintenance training for ACES staff so that staff are able to provide hardware support for all other times the Offeror is unable to provide hardware repair in a timely manner outside of the PPM. The Successful Offeror shall supply hardware maintenance procedural documentation, training, and manuals necessary to support this effort.

All proposed maintenance and support solutions shall include the following features and meet all requirements of this section.

#### 4.6.3 General Service Provisions

The Successful Offeror shall be responsible, at its own expense, for the repair or replacement of any failing hardware component that it supplies and correction of defects in software that it provides as part of the system.

At its sole discretion, ACES may request advance replacement of components which show a pattern of failures which reasonably indicates that future failures may occur in excess of reliability targets, or for which there is a systemic problem that prevents effective use of the system.

Hardware failures due to environmental changes in facility power and cooling systems which can be reasonably anticipated (such as brown-outs, voltage-spikes or cooling system failures) are the responsibility of the Offeror.

#### 4.6.4 Software and Firmware Update Service

The Successful Offeror shall provide an update service for all software and firmware provided for the duration of the Warranty plus Maintenance period. This shall include new releases of software/firmware and software/firmware patches as required for normal use. The Successful Offeror shall integrate software fixes, revisions or upgraded versions in

supplied software, including community software (e.g. Linux or Lustre), and make them available to ACES within twelve (12) months of their general availability. The Successful Offeror shall provide prompt availability of patches for cybersecurity defects.

#### 4.6.5 Call Service

The Successful Offeror shall provide contact information for technical personnel with knowledge of the proposed equipment and software. These personnel shall be available for consultation by telephone and electronic mail with ACES personnel. In the case of degraded performance, the Successful Offeror's services shall be made readily available to develop strategies for improving performance, i.e. patches, workarounds.

#### 4.6.6 On-site Parts Cache

The Successful Offeror shall maintain a parts cache on-site at the ACES facilities. The parts cache shall be sized and provisioned sufficiently to support all normal repair actions for two weeks without the need for parts refresh. The initial sizing and provisioning of the cache shall be based on Offeror's Mean Time Between Failure (MTBF) estimates for each FRU and each rack, scaled based on the number of FRU's and racks delivered. The parts cache configuration will be periodically reviewed for quantities needed to satisfy this requirement, and adjusted if necessary, based on observed FRU or node failure rates. The parts cache will be resized, at the Offeror's expense, should the on-site parts cache prove to be insufficient to sustain the actually observed FRU or node failure rates.

#### 4.6.7 On-Site Node Cache

The Successful Offeror shall also maintain an on-site spare node inventory of at least 1% of the total nodes in all of the system. These nodes shall be maintained and tested for hardware integrity and functionality utilizing the Hardware Support Cluster defined below if provided.

#### 4.6.8 Hardware Support Cluster

The Successful Offeror shall provide a Hardware Support Cluster (HSC). The HSC shall support the hot spare nodes and provide functions such as hardware burn-in, problem diagnosis, etc. The Successful Offeror shall supply sufficient racks, interconnect, networking, storage equipment and any associated hardware/software necessary to make the HSC a stand-alone system capable of running diagnostics on individual or clusters of HSC nodes. ACES will store and inventory the HSC and other on-site parts cache components.

#### 4.6.9 DOE Q-Cleared Technical Service Personnel

The Crossroads system will be installed in security areas that require a DOE Q-clearance for access. It will be possible to install the system with the assistance of uncleared US citizens or L-cleared personnel, but the Successful Offeror shall arrange and pay for appropriate 3<sup>rd</sup> party security escorts. The Successful Offeror shall obtain necessary clearances for on-site support staff to perform their duties.

## **5 Delivery and Acceptance**

Testing of the system shall proceed in three steps: pre-delivery, post-delivery, and acceptance. Each step is intended to validate the system and feeds into subsequent activities. Sample Acceptance Test plans (Appendix A) are provided as part of the Request for Proposal.

### **5.1 Pre-delivery Testing**

The ACES team and the Successful Offeror shall perform pre-delivery testing at the factory on the hardware to be delivered. Any limitations for performing the pre-delivery testing should be identified in the Offeror's proposal, including scale and licensing limitations (if any). During pre-delivery testing, the Successful Offeror shall:

- Demonstrate RAS capabilities and robustness using simple fault injection techniques, such as disconnecting cables, powering down subsystems, or installing known bad parts.
- Demonstrate functional capabilities on each segment of the system built, including the capacity to build applications, schedule jobs, and run them using a customer-provided testing framework. The root cause of application failure must be identified prior to system shipping.
- Provide a file system sufficiently provisioned to support the suite of tests.
- Provide onsite and remote access to the ACES team to monitor testing and analyze results.
- Instill confidence in the ability to conform to the statement of work.

### **5.2 Site Integration and Post-delivery Testing**

The ACES team and the Successful Offeror staff shall perform site integration and post-delivery testing on the fully delivered system. Limitations and/or special requirements may exist for access to the onsite system by the Offeror.

- During post-delivery testing, the pre-delivery tests shall be run on the full system installation.
- Where applicable, tests shall be run at full scale.

### 5.3 Acceptance Testing

The ACES team and the Successful Offeror staff shall perform onsite acceptance testing on the fully installed system. Limitations and/or special requirements may exist for access to the onsite system by the Offeror.

- 5.3.1 The Successful Offeror shall demonstrate that the delivered system conforms to the subcontract's Statement of Work.

## 6 Risk and Project Management

The Offeror shall propose a risk management strategy and project management plan for the Crossroads system that is closely coordinated between the subcontracts for Triad National Security, LLC.

- 6.1.1 The Offeror shall Propose a risk management strategy for the system in the event of technology problems or scheduling delays that affect delivery of the system or achievement of performance targets in the proposed timeframe. Offeror shall describe the impact of substitute technologies (if any) on the overall architecture and performance of the system in particular addressing the four technology areas listed below:
- Processor
  - Memory
  - High-speed interconnect
  - Platform storage
- 6.1.2 The Offeror shall identify any other high-risk areas and accompanying mitigation strategies for the system.
- 6.1.3 The Offeror shall provide a clear plan for effectively responding to software and hardware defects and system outages at each severity level and document how problems or defects will be escalated.
- 6.1.4 The Offeror shall propose a roadmap showing how their response to this Request for Proposal aligns with their plans for Exascale computing.
- 6.1.5 The Offeror shall identify additional capabilities, including:
- Its ability to produce and maintain the system for the life of the system
  - Its ability to achieve specific quality assurance, reliability, availability and serviceability goals
  - Its in-house testing and problem diagnosis capability, including hardware resources at appropriate scale

- 6.1.6 The Offeror shall provide project management specifics for the ACES team detailed as part of the Request for Proposal document. Please see Appendix B for further information.

## 7 Documentation and Training

The Successful Offeror shall provide documentation and training to effectively operate, configure, maintain, and use the systems to the ACES team and users of the Crossroads system. The ACES team may, at their option, make audio and video recordings of presentations from the Successful Offeror's speakers at public events targeted at the NNSA user communities (e.g., user training events, collaborative application events, best practices discussions, etc.). The Successful Offeror shall grant the ACES team user and distribution rights of documentation provided by the Offeror, session materials, and recorded media to be shared with other DOE Labs' staff and all authorized users and support staff for Crossroads.

### 7.1 Documentation

- 7.1.1 The Successful Offeror shall provide documentation for each delivered system describing the configuration, interconnect topology, labeling schema, hardware layout, etc. of the system as deployed before the commencement of system acceptance testing.
- 7.1.2 The Successful Offeror shall supply and support system and user-level documentation for all components before the delivery of the system. Upon request by the laboratories, the Successful Offeror shall supply additional documentation necessary for operation and maintenance of the system. All user-level documentation shall be publicly available.
- 7.1.3 The Successful Offeror shall distribute and update all documentation electronically and in a timely manner. For example, changes to the system shall be accompanied by relevant documentation. Documentation of changes and fixes may be distributed electronically in the form of release notes. Reference manuals may be updated later, but effort should be made to keep all documentation current.

### 7.2 Training

- 7.2.1 The Successful Offeror shall provide the following types of training at facilities specified by ACES:

Class Type	Number of Classes
System Operations and Advanced Administration	2



User Programming	3
------------------	---

7.2.2 The Offeror shall describe all proposed training and documentation relevant to the proposed solutions utilizing the following methods:

- Classroom training
- Onsite training
- Online documentation
- Online training

## 8 References

ACES schedule and high-level information can be found at the primary Crossroads website (<http://crossroads.lanl.gov/>).

Crossroads benchmarks and workflows whitepaper can be found at the Crossroads Benchmark and Workflows [website](#).

High Performance Computing Power Application Programming Interface Specification website (<http://powerapi.sandia.gov/>).

## **Appendix A: Sample Acceptance Plan**

### **Appendix A-1: ACES Sample Acceptance Plan**

Testing of the system shall proceed in three steps: pre-delivery, post-delivery and acceptance. Each step is intended to validate the system and feeds into subsequent activities.

#### **Pre-delivery (Factory) Test**

The Subcontractor shall demonstrate all hardware is fully functional prior to shipping. If the system is to be delivered in separate shipments, each shipment shall undergo pre-delivery testing. If the Subcontractor proposes a development system subcomponent, Triad recognizes that the development system is not part of the pre-delivery acceptance criteria.

ACES and Subcontractor staff shall perform pre-delivery testing at the factory on the hardware to be delivered. Any limitations for performing the pre-delivery testing need to be identified including scale and licensing limitations.

- Demonstrate RAS capabilities and robustness, using simple fault injection techniques such as disconnecting cables, powering down subsystems, or installing known bad parts.
- Demonstrate functional capabilities on each segment of the system built, including the capability to build applications, schedule jobs, and run them using the customer-provided testing framework. The root cause of any application failure must be identified.
- The Offeror shall provide a file system sufficiently provisioned to support the suite of tests.
- Provide onsite and remote access for ACES staff to monitor testing and analyze results.
- Instill confidence in the ability to conform to the statement of work.

#### **Pre-Delivery Assembly**

- The Subcontractor shall perform the pre-delivery test of Crossroads or agreed-upon sub-configurations of Crossroads at the Subcontractor's location prior to shipment. At its option, ACES may send a representative(s) to observe testing at the Subcontractor's facility. Work to be performed by the Subcontractor includes:
  - All hardware installation and assembly
  - Burn in of all components
  - Installation of software

- Implementation of the ACES-specific production system-configuration and programming environment
- Perform tests and benchmarks to validate functionality, performance, reliability, and quality
- Run benchmarks and demonstrate that benchmarks meet performance commitments.

#### Pre-Delivery Configuration

- TBD

#### Pre-Delivery Test

Subcontractor shall provide ACES on-site access to the system in order to verify that the system demonstrates the ability to pass acceptance criteria.

The pre-delivery test shall consist of (but is not limited to) the following tests:

<b>Name of Test</b>	<b>Pass Criteria</b>
System power up	All nodes boot successfully
System power down	All nodes shut down
Unix commands	All UNIX/Linux and vendor specific commands function correctly
Monitoring	Monitoring software shows status for all nodes
Reset	“Reset” functions on all nodes
Power On/Off	Power cycle all components of the entire system from the console
Fail Over/Resilience	Demonstrate proper operation of all fail-over or resilience mechanisms
Full Configuration Test	Pre-delivery system can efficiently run applications that use the entire compute resource of the pre-delivery system. The applications to be run will be drawn from the 72-hour test runs, scaled to the pre-delivery configuration
Benchmarks	Benchmarks shall achieve performance within the limits of pre-delivery configuration

Name of Test	Pass Criteria
72 Hour test	100% availability of the pre-delivery system for a 72-hour test period while running an agreed-upon workload that exercises at least 99% of the compute resources

## Post-delivery Integration and Test

### Post-delivery Integration

During Post-Delivery Integration, the Subcontractor's system(s) shall be delivered, installed, fully integrated, and shall undergo Subcontractor stabilization processes. Post-delivery testing shall include replication of all of the pre-delivery testing steps, along with appropriate tests at scale, on the fully integrated platform. Where applicable, tests shall be run at full scale.

### Site Integration

When the Subcontractor has declared the system to be stable, the Subcontractor shall make the system available to ACES personnel for site-specific integration and customization. Once the Subcontractor's system has undergone site-specific integration and customization, the acceptance test shall commence.

## Acceptance Test

The Acceptance Test Period shall commence when the system has been delivered, physically installed, and undergone stabilization and site-specific integration and customization completed. The duration of the Acceptance Test period is defined in the Statement of Work.

All tests shall be performed on the initial production configuration as defined by ACES.

The Subcontractor shall supply source code used, compile scripts, output, and verification files for all tests run by the Subcontractor. All such provided materials become the property of Triad.

All tests shall be performed on the initial production configuration of the Crossroads system as it will be deployed to the ACES user community. ACES may run all or any portion of these tests at any time on the system to ensure the Subcontractor's compliance with the requirements set forth in this document.

The acceptance test shall consist of a Functionality Demonstration, a System Boot Test, a System Resilience Test, a Performance Test, and an Availability Test, performed in that order.

### Functionality Demonstration

Subcontractor and ACES will perform the Functionality Demonstration on a dedicated system. The Functionality Demonstration shall show that the system is configured and functions in accordance with the statement of work. Demonstrations shall include, but are not limited to, the following:

- Remote monitoring, power control and boot capability
- Network connectivity
- File system functionality
- Batch system
- System management software
- Program building and debugging (e.g. compilers, linkers, libraries, etc.)
- Unix functions

#### System Boot Test

Subcontractor and ACES will perform the System Boot Test on a dedicated system. The System Boot Test shall show that the system is configured and functions in accordance with the statement of work. Demonstrations shall include, but are not limited to, the following:

Two successful system cold boots to production state, with no intervention to bring the system up. Production state is defined as running all system services required for production use and being able to compile and run parallel jobs on the full system. In a cold boot, all elements of the system (compute, login, I/O) are completely powered off before the boot sequence is initiated. All components are then powered on.

- Single node power-fail/reset test: Failure or reset of a single compute node shall not cause system-wide failure.

#### System Resilience Test

Subcontractor and ACES will perform the System Resilience Test on a dedicated system. The System Resilience Test shall show that the system is configured and functions in accordance with the statement of work.

All system resilience features of Crossroads shall be demonstrated via fault-injection tests when running test applications at scale. Fault injection operations should include both graceful and hard shutdowns of components. The metrics for resilience operations include correct operation, any loss of access or data, and time to complete the initial

recovery plus any time required to restore (fail-back) a normal operating mode for the failed components.

#### Performance Test

Crossroads system performance and benchmark tests are fully documented in the Statement of Work along with guidance and test information found at the Crossroads [website](#).

The Subcontractor shall run the Crossroads tests and application benchmarks, full configuration test, external network test and file system metadata test as described in the Application and Benchmark Run Rules document. Benchmark answers must be correct, and each benchmark result must meet or exceed performance commitments in the performance requirements section.

Benchmarks must be run using the supplied resource management and scheduling software. Except as required by the run rules, benchmarks need not be run concurrently. If requested by Triad, Subcontractor shall reconfigure the resource management software to utilize only a subset of compute nodes, specified by Triad.

#### JMTTI and System Availability Testing

The JMTTI and System Availability Test will commence after successful completion of the Functionality Demonstration, System Test and Performance Test. ACES will perform the JMTTI and Availability Test.

The Crossroads system must demonstrate the JMTTI and availability metrics defined in the Statement of Work, within an agreed-upon period of time. An automated job launch and outcome analysis tool, such as the Pavilion HPC Testing Framework, shall be used to manage an agreed-upon workload that will be used to measure the reliability of individual jobs. These jobs shall be a mixture of benchmarks from the Performance Test and other applications.

Every test in the JMTTI and System Availability Test workload shall obtain a correct result in both dedicated and non-dedicated modes:

- In dedicated mode, each benchmark in the Performance Test shall meet the performance commitment specified in the Statement of Work. In non-dedicated mode, the mean performance of each performance test shall meet or exceed the performance commitment specified in the Statement of Work
- During the JMTTI and System Availability Test, ACES shall have full access to the system and shall monitor the system. Triad and users designated by Triad shall submit jobs through the Crossroads resource management system.

- During the JMTTI and System Availability Test, the Subcontractor shall adhere to the following requirements:
  - All hardware and software shall be fully functional at the end of the JMTTI and Availability Test. Any down time required to repair failed hardware or software shall be considered an outage unless it can be repaired without impacting system availability.
  - Hardware and software upgrades shall not be permitted during the last 7 days of the JMTTI and Availability Test. The system shall be considered down for the time required to perform any upgrades, including rolling upgrades.
  - No significant (i.e. levels 1, 2 or 3) problems shall be open during the last 7 days.
- During the JMTTI and Availability Testing period, if any system software upgrade or significant hardware repairs are applied, the Subcontractor shall be required to run the Performance Tests and demonstrate that the changes incur no loss of performance. At its option, Triad may also run any test deemed necessary. Time taken to run the Performance and other tests shall not count as downtime, provided that all tests perform to specifications.

#### Definitions for Node and System Failures

The baseline of interrupts, as used in the JMTTI and SMTBI calculations, shall include, but may not be limited to, the following circumstances:

- A node shall be defined as down if a hardware problem causes Subcontractor supplied software to crash or the node is unavailable. Failures that are transparent to Subcontractor-supplied software because of redundant hardware shall not be classified as a node being down as long as the failure does not impact node or system performance. Low severity software bugs and suggestions (e.g. wrong error message) associated with Subcontractor supplied software will not be classified as a node being down.
- A node shall be classified as down if a defect in the Subcontractor supplied software causes a node to be unavailable. Communication network failures external to the system, and user application program bugs that do not impact other users shall not constitute a node being down.
- Repeat failures within eight hours of the previous failure shall be counted as one continuous failure.

- The Subcontractor's system shall be classified as down (and all nodes shall be considered down) if any of the following requirements cannot be met ("system-wide failures"):
  - Complete a POSIX `stat` operation on any file within all Subcontractor-provided file systems and access all data blocks associated with these files.
  - Complete a successful interactive login to the Subcontractor's system. Failures in the ACES network do not constitute a system-wide failure.
  - Successfully run any part of the performance test. The Performance Test consists of the Crossroads Benchmarks, the Full Configuration Test and the External Network Test.
  - Full switch bandwidth is available. Failure of a switch adapter in a node does not constitute a system-wide failure. However, failure of a switch would constitute failure, even if alternate switch paths were available, because full bandwidth would not be available for multiple nodes.
  - User applications can be launched and/or completed via the scheduler.
- Other failures in Subcontractor supplied products and services that disrupt work on a significant portion of the nodes shall constitute a system-wide outage.
- If there is a system-wide outage, Triad shall turn over the system to the Subcontractor for service when the Subcontractor indicates they are ready to begin work on the system. All nodes are considered down during a system-wide outage.
- Downtime for any outage shall begin when Triad notifies the Subcontractor of a problem (e.g. an official problem report is opened) and, for system outages, when the system is made available to the Subcontractor. Downtime shall end when:
  - For problems that can be addressed by bringing up a spare node or by rebooting the down node, the downtime shall end when a spare node or the down node is available for production use.
  - For problems requiring the Subcontractor to repair a failed hardware component, the downtime shall end when the failed component is returned to Triad and available for production use.

For software downtime, the downtime shall end when the Subcontractor supplies a fix that rectifies the problem or when Triad reverts to a prior copy of the failing software that does not exhibit the same problem. A failure due



to ACES or to other causes out of the Subcontractor's control shall not be counted against the Subcontractor unless the failure demonstrates a defect in the system. If there are any disagreements as to whether a failure is the fault of the Subcontractor or ACES, they shall be resolved prior to the end of the acceptance period.

## **Appendix B: Triad Specific Project Management Requirements**

### **Appendix B-1: Triad Project Management Requirements**

#### **NOTE:**

The following requirements apply to the project management of the delivery of the system proposed by the subcontractor.

#### **Project Management**

The development, pre-shipment testing, installation and acceptance testing of the Crossroads system is a complex endeavor and will require close cooperation between the Subcontractor, Triad National Security, LLC (Triad, TNS), and ACES. There shall be quarterly executive reviews by corporate officers of the Subcontractor, ACES, and representatives of DOE NNSA, to assess the progress of the project.

#### **Project Planning Workshop**

- Triad and Subcontractor shall schedule and complete a workshop to mutually understand and agree upon project management goals, techniques, and processes.
- The workshop shall take place no later than award + 45 days

#### **Project Plan**

- Delivery Milestone: no later than award + 60 days

Subcontractor shall provide Triad with a detailed Project Plan – which includes a detailed Work Breakdown Structure (WBS). The Project Plan shall contain all aspects of the proposed Subcontractor's solution and associated engineering (hardware and software) and support activities.

The Project Plan shall address or include:

- Program Management
- High Assurance Delivery Process

#### **WBS:**

- Facilities Planning (e.g., floor, power & cooling, cabling);
- Computer Hardware Planning;
- Installation & Test Planning;
- Deployment and Integration Milestones
- System Stability Planning;
- System Scalability Planning;

- Software Plan
- Testing
- Development
- Testing
- Deployment
- Risk Assessment & Risk Mitigation
- Staffing;
- On-site Warranty and Maintenance and Support Planning;
- Training & Education;

**Project Plan – Program Management**

At a minimum, the Project Plan – Program Management Section shall:

- Identify, by name, the Program Management Team members;
  - Identify, by name, the lead Crossroads System Architect
  - Identify, by name, the Crossroads System RAS Point of Contact
  - Describe the roles and responsibilities of the Team members;
  - List Subcontractor's Management Contacts;
  - Define and institutionalize the Periodic Progress Review process with regard to frequency (daily, weekly, monthly, quarterly, and annually), level (support, technical, and executive), and escalation procedures.
- Additionally, the Project Plan – Program Management Section shall detail the joint activities of the Subcontractor and Triad to monitor and assess the overall Program Performance.
  - Triad will furnish the Subcontractor with a top-10 list of problems and issues. The Subcontractor is responsible for appointing a point of contact for each of the items on the list. This list shall be reviewed weekly.
  - All Subcontractor Program Management shall interface with the designated Triad Crossroads project manager.
  - The WBS will be updated by the Subcontractor monthly and reviewed for approval by Triad
  - The Subcontractor Project Plan shall be updated by the Subcontractor quarterly and reviewed for approval by Triad

**Project Plan - High Assurance Hardware Delivery Process**

Subcontractor shall provide Triad with a high assurance delivery process and certification program for hardware deliverables of all stages of the deployment and operational use by the ASC Applications Community of the systems.

All assets delivered shall be, at a minimum, factory-tested and field-certified;

A “pre-delivery test” shall take place at the factory prior to each shipment. Functional diagnostics and agreed upon ACES applications shall be executed to verify the proper functioning of each system prior to shipment. Problems identified as a result of these tests shall be corrected prior to shipment. Assets that have successfully completed this pre-delivery test are “pre-verified.”

### **Project Plan - High Assurance Software Delivery Process**

Subcontractor shall provide Triad with a high assurance delivery process and certification program for software deliverables of all stages of the deployment and operational use by the NNSA ASC tri-lab simulation community of the Crossroads systems. In addition, Subcontractor shall provide Triad with documentation of Subcontractor’s anticipated software release schedules during lifetime of the subcontract. This includes major and minor releases, updates, and fixes as well as expected beta-level availability.

- While Beta software and/or pre-GA software is anticipated to be installed and run on these systems, however all such installations are subject to Triad approval;
- Subcontractor shall provide Triad with a list of interdependencies between hardware and software as they pertain to the delivered systems;

### **Project Plan – WBS, Milestones**

Subcontractor shall define appropriate high-level Milestones for the execution of the delivery and acceptance of the Crossroads system.

### **Project Plan – WBS, Facilities Planning**

Compliant with the requirements of the Facilities described in the Technical Requirements.

### **Project Plan – WBS, System Stability Planning**

Scalable systems of the size being delivered can at times prove difficult to predict in terms of stability. The number of components can have a significant effect on the stability and may provide some scalability problems in terms of stability of the system. Triad requires a plan to progressively qualify a series of configurations of increasing complexity, in terms of both processor counts and interconnect topology.

Subcontractor shall be responsible for delivering a Stabilization Plan that includes the following:

- Plan objectives
- Target Goals for Stability, as agreed to jointly with Triad and ACES
- Technical Strategy
- Roles and responsibilities
- Testing Plan
- Progress Evaluation Checkpoints
- Contingencies

**Project Plan – Staffing:**

- Staff Support shall be for the life of the subcontract.
- Subcontractor shall identify its members of the Project Team.

**Project Plan – On-site Warranty and Maintenance and Support Planning**

- On-site Warranty and Maintenance and Support shall be for the life of the subcontract
- On-site Warranty and Maintenance and Support shall include Subcontractor's preventive maintenance schedule.
- On-site Warranty and Maintenance and Support shall include logging and weekly reporting of all interruptions to service. At a minimum, the Subcontractor shall enter all interrupt logging into Triad tracking system.

**Project Plan – Training and Education**

- In addition to Subcontractor's usual and customary customer Training and Education program, Subcontractor shall allow ACES staff access to Subcontractor's internal Training & Education program;
- Training and Education Support shall be for the life of the subcontract.

**Project Plan – Risk Assessment and Risk Mitigation**

- Subcontractor shall provide Triad with a Risk Management Plan that identifies and addresses all identified risks.
- Subcontractor shall provide a risk management strategy for the proposed system in case of technology problems or scheduling delays that affect availability or achievement of performance targets in the proposed timeframe. Subcontractor shall describe the impact of substitute technologies on the overall architecture and performance of the system. In particular, the subcontractor shall address the technology areas listed below:

- Processor
- Memory
- High-Speed Interconnect
- Platform Storage and all other I/O subsystems
- Subcontractor shall continuously monitor and assess the risks involved for those major technology components that Subcontractor identifies to be on the Critical Path (i.e., Risk Assessment);
- Subcontractor shall provide Triad with timely and regular updates regarding Subcontractor's Risk Assessment;
- Subcontractor shall provide Triad with a Risk Mitigation Plan. Each risk mitigation strategy shall be subject to Triad approval. Such Risk Mitigation Plan shall include:
  - Risks Categorization – Risks shall be categorized according to
  - Probability of occurrence (Low, medium, or high)
  - Impact to the program if they occur (low, medium, or high)
  - Dates for Risk Mitigation Decision Points Identified
  - Execution of mitigation plans are subject to Triad approval and may include:
    - *Technology Substitution* – subject to the condition that substituted technologies shall not have aggregate performance, capability, or capacity less than originally proposed;
    - *3rd Party Assistance* – especially in areas of critical software development;
    - *Source Code Availability* – especially in the areas of Operating Systems, Communication Libraries;
    - *Performance Compensation* – possibility of compensating for performance shortfalls via additional deliveries.
  - Subcontractor's Risk Mitigation Plan will be reviewed quarterly by Triad.

## Definitions and Glossary

**Baseline Memory:** High performance memory technologies such as DDR-DRAM, HBM, and HMC, for example, that may be included in the systems memory capacity requirement. It does not include memory associated with caches.

**Coefficient of Variation:** The ratio of the standard deviation to the mean.

**Coldboot:** Full power-on of a system from a non-energized state, such as a post power outage situation. It can be assumed that facilities power, water, network, and site infrastructure services have been returned to service and timings for all offered file systems and cluster can begin from this point.

**Delta-Ckpt:** The time to checkpoint 80% of aggregate memory of the system to persistent storage. For example, if the aggregate memory of the compute partition is 3 PiB, Delta-Ckpt is the time to checkpoint 2.4 PiB. Rationale: This will provide a checkpoint efficiency of about 90% for full system jobs.

**Ejection Bandwidth:** Bandwidth leaving the node (i.e., NIC to router).

**Full Scale:** All of the compute nodes in the system. This may or may not include all available compute resources on a node, depending on the use case.

**Idle Power:** The projected power consumed on the system when the system is in an **Idle State**.

**Idle State:** A state when the system is prepared to but not currently executing jobs. There may be multiple idle states.

**Injection Bandwidth:** Bandwidth entering the node (i.e., router to NIC).

**Job Interrupt:** Any system event that causes a job to unintentionally terminate.

**Job Mean Time to Interrupt (JMTTI):** Average time between job interrupts over a given time interval on the full scale of the system. Automatic restarts do not mitigate a job interrupt for this metric.

**JMTTI/Delta-Ckpt:** Ratio of the JMTTI to Delta-Ckpt, which provides a measure of how much useful work can be achieved on the system.

**Nameplate Power:** The maximum theoretical power the system could consume. This is a design limit, likely not achievable in operation, commonly specified on electrical equipment labels and used for power provisioning design per National Electrical Code (NEC, NFPA 70).

**Nominal Power:** The projected power consumed on the system by the ACES workflows (e.g., a combination of the ACES benchmark codes running large problems on the entire system).

**Operational Capability:** Real, usable capabilities in production operation, not theoretical capabilities.

**Peak Power:** The projected power consumed by an application that utilizes the maximum achievable power consumption such as DGEMM.

**Platform Storage:** Any nonvolatile storage that is directly usable by the system, its system software, and applications. Examples would include disk drives, RAID devices, and solid-state drives, no matter the method of attachment.

**Rolling Upgrades/Rolling Rollbacks:** A rolling upgrade or a rollback is defined as changing the operating software or firmware of a system component in such a way that the change does not require synchronization across the entire system. Rolling upgrades and rollbacks are designed to be performed with those parts of the system that are not being worked on remaining in full operational capacity.

**System Interrupt:** Any system event, or accumulation of system events over time, resulting in more than 1% of the compute resource being unavailable at any given time. Loss of access to any dependent subsystem (e.g., platform storage or service partition resource) will also incur a system interrupt.

**System Mean Time Between Interrupt (SMTBI):** Average time between system interrupts over a given time interval.

**System Availability:** 
$$\frac{(\text{time in period} - \text{time unavailable due to outages in period})}{(\text{time in period} - \text{time unavailable due to scheduled outages in period})} * 100$$

**System Initialization:** The time to bring 99% of the compute resource and 100% of any service resource to the point where a job can be successfully launched.

**Warmboot:** The cluster/file system management servers being booted and configured.